

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 January 2001 (04.01.2001)

PCT

(10) International Publication Number
WO 01/01675 A2

(51) International Patent Classification⁷: H04N 5/217

(21) International Application Number: PCT/US00/18046

(22) International Filing Date: 29 June 2000 (29.06.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/345,167 30 June 1999 (30.06.1999) US
09/343,934 30 June 1999 (30.06.1999) US
09/464,364 15 December 1999 (15.12.1999) US
09/602,547 23 June 2000 (23.06.2000) US

(71) Applicant: LOGITECH, INC. [US/US]; 6505 Kaiser Drive, Fremont, CA 94555-3615 (US).

(72) Inventors: HSU, Mark; 101 Blaisdell Way, Fremont, CA 94536 (US). NORCROSS, Mitchell; 43511 Ocaso Corte, Fremont, CA 94539 (US). AUBERGER, Georges; 615

Palm Avenue, Los Altos, CA 94022 (US). ZIMMERMANN, Remy; 1549 Molitor Road, Belmont, CA 94022 (US). MAGGI, Sergio; 2144 Balmoral Court, San Mateo, CA 94413 (US). SANCHEZ, George; 657 Los Pinos Place, Fremont, CA 94539 (US). BILLERBECK, Bryed; 441 Martens Avenue, Mountain View, CA 94040 (US). LI, Wei; 18820 Tilson Avenue, Cupertino, CA 95014 (US). LABROUSSE, Junien; 3750 Corina Way, Palo Alto, CA 94303 (US).

(74) Agents: HAUGHEY, Paul, C. et al.; Townsend And Townsend And Crew LLP, Two Embarcadero Center, 8th Floor, San Francisco, CA 94111-3834 (US).

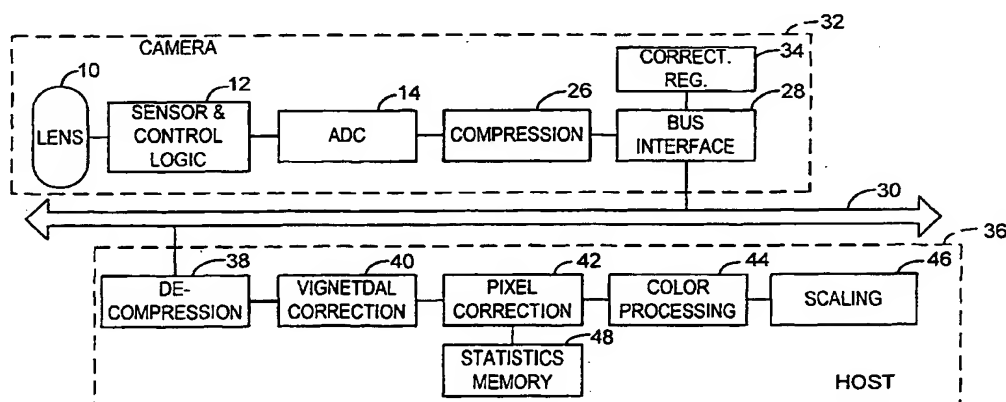
(81) Designated States (*national*): CN (utility model), DE (utility model), KR.

Published:

— Without international search report and to be republished upon receipt of that report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: VIDEO CAMERA WITH MAJOR FUNCTIONS IMPLEMENTED IN HOST SOFTWARE



(57) Abstract: A low cost camera by implementing the major functions in host software is provided. This is accomplished by sending raw, digitized data from the camera directly to the host. The increased volume of raw data is handled by either an improved compression/decompression scheme using lossless compression, using lossy compression or using a shared bus with higher bandwidth. By moving such functions as colour processing and scaling to the host, the pixel correction can also be moved to the host. This in turn allows the elimination of the frame buffer memory from the camera. Finally, the camera can use a low cost lens by implementing vignetting, distortion, gamma or aliasing correction with a correction value stored in a register of the camera for later access by the host to perform corrections.

VIDEO CAMERA WITH MAJOR FUNCTIONS IMPLEMENTED IN HOST SOFTWARE

CROSS-REFERENCES TO RELATED APPLICATIONS

5 This application is a continuation-in-part of Application No. 09/345,167,
filed June 30, 1999, entitled "IMAGE SENSOR BASED VIGNETTING
CORRECTION", Application No. 09/464,364, filed December 15, 1999, entitled
"DYNAMIC ANOMALOUS PIXEL DETECTION AND CORRECTION", and
Application No. 09/343,934, filed June 30, 1999, entitled "FAST DECODING". All of
10 which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

The present invention relates to video cameras, and in particular to video
cameras connected to a computer or other intelligent device by a shared interface, such as
15 the universal serial bus (USB).

Video cameras for providing both still pictures and motion pictures have
been developed which can connect to a computer to provide a digitized image. Such
cameras can be connected to the computer via a shared bus, such as the USB. This limits
the bandwidth of data that can be sent, requiring compression of the video data.
20 Typically, cameras will do some of the preprocessing of the image in the camera
electronics, thus both off-loading processing from the CPU of the computer and
potentially reducing the amount of data that has to be compressed and sent over the bus.

With such a design, some functions are required to be done in the camera.
For example, a CCD or CMOS sensor in the camera may have defective detector
25 positions which do not properly collect light. These can be compensated for by averaging
the amount of light detected at neighboring detector sites and substituting this for the
value provided by the defective site. This type of processing needs to be done before
other types of processing which are now typically done in a video camera, such as color
processing and scaling.

30 Fig. 1 is a block diagram of a typical video camera which can connect over
a bus to a host. A lens 10 focuses light onto a sensor 12 with associated control logic for
reading out the sensor data. The data is provided to an analog-to-digital converter (ADC)
14, where it is digitized. The data is then stored in a frame buffer memory 16. Pixel

correction can be applied to the data with a circuit 18. Correction data may be stored in a programmable memory 20. This memory may be programmed during the manufacturing process, before the camera is even shipped.

A number of other functions are typically performed in the camera, shown in Fig. 1 in a particular order, although they may be in other orders or connected via a bus to vary when they are invoked. For example, a color processing circuit 22 may perform a number of color functions, such as converting the received data from one color format, such as YUV, to another format, such as RGB. A scaler function 24 can scale the image to fit the display desired at the host. A compression circuit 26 will compress the data to reduce the amount of data to be transferred over the shared bus. Finally, a bus interface circuit 28 implements the protocols of a bus 30 for transferring data to a remote host.

Huffman coding is one of the algorithms most widely used to compress video data. Huffman coding has been used in various image and video compression standards such as the standards specified by the Joint Photographic Experts Group (JPEG), Motion Picture Experts Group (MPEG), MPEG-II, H.261, H.263, and H.323. These standards are included herein by reference, in their entirety and for all purposes. Huffman coding is a lossless entropy coding technique. Huffman coding utilizes probability to select the shortest code word for the data that occur most frequently. For example, when encoding English text, "E" which is a very popular character in the English language can be represented by a 2-bit code word, whereas "A" which is not used as frequently can be represented by an 8-bit code word.

As a result, Huffman coding compresses data by assigning shorter code words to more frequently occurring data and longer code words to less frequently occurring data. The assigned code words can be maintained in a table which is then used by both recipients and senders. Further details regarding Huffman encoding can be found in "A Method for the Construction of Minimum Redundancy Codes," *Proceedings of the IRE*, Vol. 40, No. 9, September 1952, pp. 1098-1101, by D. A. Huffman, which is incorporated herein by reference in its entirety and for all purposes.

Moreover, Huffman code words are constructed in such a way that no code word is the prefix of another code word. This assures the unambiguous decoding of each symbol from a bitstream of data despite the variable lengths of the codes. Additional description of the characteristics of Huffman codes can be found in *Data Compression Book*, M&T Books, 1996, by M. Nelson and J. L. Gailly, which is also incorporated herein by reference in its entirety and for all purposes. On pages 31-35, this reference

proposes a decoding algorithm based on tree-tracing. This algorithm is, however, not suitable for fast decompression using most of the currently available microprocessors such as Intel[®]'s Pentium[®] family.

The vignetting effect is when the brightness around the edges of a screen is less than the brightness in the center of the screen, creating a halo effect. Professional photographers will sometimes try to intentionally create this effect with an appropriate lens or screen, to produce a softened image at the edges of a picture. However, when this effect appears unintentionally in digital photographs, it is annoying. A variety of techniques have been developed to measure and address vignetting effects.

Patent No. 4,884,140 shows an analog circuit for providing vignetting compensation for a video camera using a zoom lens which causes vignetting.

Patent No. 5,434,902 shows measuring the vignetting effect for an x-ray examination apparatus by using a constant brightness image. A correction factor is then stored in the memory for each pixel.

Patent No. 5,576,797 shows the detection of vignetting effects in a camera with a focus detecting device.

Patent No. 5,381,174 shows the correcting of vignetting due to operation of a zoom lens by using a field frequency sawtooth signal.

Patent No. 4,816,663 shows the detection of vignetting between a photo taking lens and focus detecting optical system.

Digital cameras for use with a personal computer for personal teleconferencing have become cheaper and cheaper. This puts pressure on the camera manufacturers to use cheaper lenses, which have more of a vignetting effect. In addition, the price pressure forces manufacturers to use fewer semiconductor chips, thus making the use of semiconductor memory to store vignetting corrections undesirable. Accordingly, there is a need for a digital camera which is inexpensive to manufacture and also corrects for vignetting of an inexpensive lens.

The processes used to produce image sensors, though highly advanced, are not perfect. Photosites in an image sensor array can vary in their sensitivity or ability to convert light into electric currents. Some variation is expected and tolerable, but on occasion, individual photosites in an array can be dead, inactive or over active. This can result in a displayed image having a dark or black dot, bright white dot or incorrect color dot, all of which are unwanted and are undesirable to the end user. Those in the digital imaging field have also recognized the problems caused by improper digitization of target

images and have proposed various solutions to detect and correct defective pixels. These prior art methods and systems can be broadly classified into two groups.

On one hand, there are various hardware-based methods for detecting and correcting defective pixels taken from an image sensor array. These hardware-based methods and systems are quite common to video camera manufacturers. Most conventional video cameras using a solid state image pickup device, incorporate a defective pixel detection and correction circuit for correcting defective pixels taken from an image sensor array. The defective pixels are produced when or after the image sensor array is manufactured. In such a camera, a defective pixel correction circuit detects a defective pixel and stores the position data and various data related to the defective pixel in a read only memory (ROM) or the like. Then, when the video camera is in use, pixel data from the defective pixel is replaced by data from a pixel near the defective pixel. One such method and system is disclosed in U.S. Pat. No. 5,796,430. The disadvantage of such approaches is the need to incorporate and program memory devices during product assembly and testing which add costs and delays. Additionally, the hardware device costs will also increase due to the need to add the correcting circuitry or logic to the video processing application specific integrated circuit (ASIC).

On the other hand, there are various software-based methods for detecting and correcting defective pixels in digital images. One such method and system is disclosed in U.S. Pat. No. 5,982,946. Such software-based methods are generally aimed at correcting bad image pixels in an already digitized and fully processed image. These already digitized images are in their final displayable and stored forms which have already been through color processing, compression and various other processing that are performed on the raw data which is read off an image sensor array. Therefore, such software-based methods which treat completely processed digital images, may correct anomalous pixels that are caused by any number of artifacts such as dust or dirt on the original scene, dust or dirt on the lens which was used to capture the scene as well as anomalous pixels in the digitized image which were caused by defective photosites. These methods typically rely on rather sophisticated and expensive systems and computer software to detect and correct bad pixels. These methods generally require user input to detect the location of the potentially bad image pixels. Once the user has visually scanned an entire image and flagged potentially bad image pixels, the system incorporating the software program takes over to correct the flagged defective pixels.

Besides requiring user input, such approaches are also expensive, tedious and very time consuming.

The proliferation of inexpensive PC-interfacing digital still and video camera devices requires rapid, dynamic, inexpensive and intelligent defective pixel
5 detection and correction solutions.

SUMMARY OF THE INVENTION

The present invention provides a low cost camera by implementing the major functions in host software. This is accomplished by sending raw, digitized data
10 from the camera directly to the host. The increased volume of raw data is handled by either an improved compression/decompression scheme using lossless compression, using lossy compression or using a shared bus with higher bandwidth. By moving such functions as color processing and scaling to the host, the pixel correction can also be moved to the host. This in turn allows the elimination of the frame buffer memory from
15 the camera. Finally, the camera can use a low cost lens by implementing vignetting correction with a vignetting correction value stored in a register of the camera for later access by the host to perform corrections.

In one embodiment, the host decompresses the transmitted data by using a processor with the capability of simultaneous operations on multiple packed pixel values,
20 such as the Intel MMX™ technology. This maintains a sufficient decompression speed for a larger amount of data with minimal impact on the frame rate. In one embodiment, a group of bits from the data stream are duplicated and provided to multiple positions in a register, where they can be simultaneously compared to multiple maximum values. This allows a quick determination of how many bits of the variable bit encoding correspond to
25 a pixel value.

In an embodiment using lossy compression, vignetting, gamma, distortion or aliasing correction and pixel correction are performed in the camera itself, since the ability to correct would be degraded by the loss of information during the lossy compression/decompression process. However, the color processing, scaling and other
30 operations are still performed in the host, achieving a low cost camera which does not require a frame buffer memory and color processing and scaling circuitry.

In one embodiment, the camera can be made low-cost by using a low-cost lens even though it has vignetting distortion. This is accomplished by providing a memory element, such as a register, in the camera. This memory element is programmed

at the time of manufacture with a value corresponding to the amount of vignetting or correction required. The register can then be read by the host during operation to determine the amount of correction required in a vignetting correction algorithm executed in the host on the received data prior to any other processing. In an alternate
5 embodiment, the memory element can store a correction or value factor for other defects in the lens or other aspects of the camera. For example, an indication of bad pixel sites could be stored for subsequent reading and correction by the host.

For a further understanding of the nature and advantages of the invention, reference should be made to the following description taken in conjunction with the
10 accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a prior art video camera.

Fig. 2 is a block diagram of one embodiment of the invention with major
15 functions moved to a host.

Fig. 3 is a diagram illustrating the operation of Huffman decoding operations in parallel.

Fig. 4 illustrates an example of a computer system used to execute methods of embodiments of the present invention;

Fig. 5 illustrates a simplified system block diagram of a typical computer
20 system 100 used to execute the methods of embodiments of the present invention;

Fig. 6 is a simplified block diagram of a SIMD system 300 in accordance with an embodiment of the present invention;

Fig. 7 illustrates a simplified block diagram of a system 400 in accordance
25 with an embodiment of the present invention;

Fig. 8 illustrates a sequential decoding method 500 for decoding Huffman encoded data in accordance with an embodiment of the present invention; and

Fig. 9 illustrates a parallel decoding method 600 for decoding Huffman encoded data in accordance with another embodiment of the present invention.

Fig. 10 is a drawing of a pixel array and brightness curve illustrating the
30 vignetting effect.

Fig. 11 is a block diagram of one embodiment of vignetting correction according to the invention.

Fig. 12 illustrates a simplified flow chart of a method of detecting and correcting defective pixels according to one embodiment of the present invention.

Fig. 13 is a block diagram of a video camera system according to one embodiment of the present invention.

5 Fig. 14 is timing diagram of the pulses provided to a CMOS sensor array in accordance with one embodiment of the present invention.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

I. OVERALL SYSTEM.

10 Fig. 2 is a block diagram of one embodiment of the invention. A camera 32 includes a lens 10, a sensor and control logic 12, and an ADC 14, as in the prior art. However, the other major functions are removed, except for a compression block 26 and bus interface 28. In addition, a correction register 34 is added to store a value corresponding to the vignetting of lens 10.

15 The camera connects over a shared bus 30 to host 36. In host 36, the blocks shown are programming blocks executed by the processor of host 36. These are a decompression block 38, a vignetting correction block 40, a pixel correction block 42, a color processing block 44 and a scaling block 46. Also shown is a statistics memory 48, which can be a portion of the host memory, for storing statistics information on pixels
20 needing correction.

II. HUFFMAN DECOMPRESSION.

Preferably, the processor used by host 36 includes the ability to perform operations on multiple packed pixels in a single register in parallel. For example, the
25 Intel MMX™ technology provides a number of instructions for such operations on packed pixels in the microprocessor registers. Accordingly, it is desirable to take advantage of such functions to increase the speed of processing. However, certain operations do not lend themselves to such simultaneous operation. For example, Huffman compression is often used for video data, and is used in the embodiment of this
30 invention. This is a variable length code compression mechanism. In other words, the number of bits corresponding to a pixel value could be 1, 2, 3, 4, etc., and can vary from pixel to pixel. Without knowing the number of bits ahead of time, and with different sizes, it is not immediately obvious how to utilize a packed pixel register which uses the same size for each of the pixel representations.

At the limited bandwidth of today's USB technology, it is desirable to be able to somehow do parallel decompression processing if the camera is to send raw data, which will have a much larger amount of data than data which has been preprocessed. This is accomplished by uniquely using the packed pixel configuration to perform the Huffman decoding.

Fig. 3 is a diagram illustrating in general how the parallel Huffman decoding proceeds. An incoming data stream 50 is illustrated by a number of letters. Each letter is an arbitrary symbol indicating a different number of bits. In the example shown, the first of number of bits, A, is duplicated four times and placed into four positions in a first register 52. A second register 54 stores four different masks. The masks enable a comparison with different portions of the bits in bit segment A. For example, the number of bits of A examined could be 1, 2, 3 and 4, respectively. These would then be compared to the maximum values for the Huffman codes for 1, 2, 3 and 4 bits. These four different maximum values are stored in another register 56. Thus, instead of sequentially comparing the number of bits A to four different values, this can be done in parallel, even though the number of bits to be compared is variable. This variation is overcome by using the masking operation. Thus, the same number of bits are stored in each pixel location in the register, but a different number of bits are compared for each of the pixel locations by virtue of the mask operation.

In the example shown, the bit segment A corresponds to being within the maximum segment M3, which is then used in an equation to index into a combined Huffman value table 58. By using the arithmetic result of the prior comparison as an index, what used to require three different table lookups in the prior art can now be done in a single table lookup, thus saving further processing time.

Fig. 4 illustrates an example of a computer system used to execute the software of the present invention. Fig. 4 shows a computer system 100 which includes a monitor 104, screen 102, cabinet 108, keyboard 214 (see Fig. 5), and mouse 110. The mouse 110 can have one or more buttons such as mouse buttons 112. The cabinet 108 can house a CD-ROM drive 106 and a hard drive (not shown) which can be utilized to store and retrieve software programs incorporating the present invention. Although the CD-ROM 106 is shown as removable media, other removable tangible media including floppy disks, tape drives, ZIP® drives, and flash memory can be utilized. The cabinet 108 can also house familiar computer components (not shown) such as a processor, memory, and the like.

Fig. 5 illustrates a simplified system block diagram of a typical computer system 100 used to execute the software of embodiments of the present invention. As shown in Fig. 1, the computer system 100 can include the monitor 104. The computer system 100 can further include subsystems such as I/O controller 204, system memory 206, central processor 208, speaker 210, removable disk 212, keyboard 214, fixed disk 216, and network interface 218. Other computer systems suitable for use with the present invention may include additional or fewer subsystems. For example, another computer system could include more than one processor 208 (i.e., a multi-processor system) or a cache memory. Arrows such as 220 represent system bus architecture of the computer system 100. However, these arrows 220 are illustrative of any interconnection scheme serving to link the subsystems.

For example, a local bus could be utilized to connect the central processor 208 to the system memory 206. Also, an image capture device such as a charge-coupled device (CCD) camera can be connected to the computer system 100 for capturing image data. The image capture device can be connected to the computer system 100 via the same or another bus architecture such as a Universal Serial Bus (USB) and the like. USB can provide plug and play support for more than 100 connected peripherals by using an identification number which is matched against a database of device drivers. The USB can also be connected to the computer system 100 through the I/O controller 204 or the network interface 218. Additionally, the computer system 100 can be configured to communicate with the Internet via, for example, the I/O controller 204 or the network interface 218. Accordingly, data can be transmitted to and from the computer system 100 by a variety of devices. The computer system 100 shown in Fig. 4 is but an example of a computer system suitable for use with the present invention. Other configurations of subsystems suitable for use with the present invention will be readily apparent to one of ordinary skill in the art.

The present invention can be embodied in any microprocessor capable of single instruction multiple data (SIMD) execution. For example, the Intel® MMX™ microprocessors, the Sun® UltraSPARC® with Visual Instructions Set, Advance Micro Device, Inc.'s® 3DNow!™, Intel® MMX™ with streaming SIMD extensions, and the like can be utilized. For the sake of example, an implementation using the Intel® MMX™ will be described. The MMX™ technology was formally introduced in January 1997 to accelerate the CPU demanding multimedia applications. The MMX™ technology provides a set of instructions (57 instructions) introduced by Intel® Corporation. It can

perform addition, subtraction, multiplication, logic, and arithmetic shifts on a unit of 64 bits in the same instruction. The unit of 64 bits can also be treated as 8 bytes, 4 words, 2 double words, or 1 quad word. Instructions for saturation arithmetic and packing/unpacking data are provided as well as those which transfer data between MMX™ registers and integer registers. Further information regarding MMX™ technology can be found in "Introduction to the Intel® Architecture MMX™ Technology Developer's Manual," Intel® (1996), which is incorporated herein by reference in its entirety and for all purposes.

Fig. 6 is a simplified block diagram of a SIMD system 300. The system 300 includes an MMX™ processor 302 which presently supports a set of 57 instructions. The MMX™ processor can be implemented within a CPU or can be alternatively implemented in a separate chip. The MMX™ processor 302 receives data from registers 304a-h. Each of the registers 304a-h are 64 bits wide. Each of these registers can hold 8x8, 16x4, 32x2, 64x1 bits of data. The MMX™ processor 302 receives the data from the registers 304a-h and performs the required operations on the data in one cycle. The processed data is then either provided to a memory 306 or back to registers 304a-h. The memory 306 can be implemented within the MMX™ processor 302 or it can alternatively be shared memory. As a result, the MMX™ processor 302 can perform the same operations on different data at the same time, hence it has an SIMD architecture.

Fig. 7 illustrates a simplified block diagram of a system 400 using the USB. A camera 402 captures images and provides them to a bitstream conversion block 404. The camera 402 can include circuitry for conversion of analog data to digital data. The camera 402 can also include circuitry to perform data compression and/or encoding. The camera 402 and the bitstream conversion block 404 are connected via connector 406. The connector 406 can be selected from a number of connectors or bus architectures which are well known to those with ordinary skill in the art. The connector 406 is preferably a USB connection. In some implementations, USB provides a 12 Mb/sec bandwidth. Another choice for connector 406 can be Firewire (IEEE 1394) which can provide a bandwidth of 100 Mb/sec, 200, 400, or 800 Mb/sec. A USB connection is preferred for some embodiments because it is a powered bus which currently provides up to 500 mA and 5 V. The power provided by USB can be utilized to run the devices connected to the bus, such as the camera 402. Additionally, USB is less costly to implement. In some embodiments, since USB does not have the same bandwidth as

Firewire, USB can be utilized where the data is first compressed prior to transmission on the USB.

Accordingly, the camera 402 can include compression circuitry to compress the captured images before sending the data to the bitstream conversion block 404. The camera 402 can be any number of devices for capturing images including a
5 CCD, complementary metal oxide semiconductor (CMOS), and the like. The bitstream conversion block 404 can be configured to convert serially transmitted data into packets of data. For example, the bitstream conversion block 404 can accumulate data for each image frame and send the accumulated data to a decompression block 408. The frame of
10 data can be any size, but is preferably 352x288 pixels. The frame can also be a block of 320x240, 176x144, or 160x120 pixels. In some embodiments, the frame of data can be a block of 640x480 pixels. The decompression block 408 decompresses and/or decodes the data received from the bitstream conversion block 404. The decoding can be in
accordance with Huffman coding, arithmetic coding, other types of entropy coding, and
15 the like. The bitstream conversion block 404 can also include buffers for storage of the data received from the camera 402 and the data sent to the decompression block 408.

The decompressed data from the decompression block 408 is then provided to a color conversion block 410. The data from the decompression block 408 can be in any format but is preferably in YUV format, where Y is luminance, U is
20 chrominance red (also known as CR), and V is chrominance blue (also known as CB). The conversion block 410 converts the YUV format data to a format suitable for a display 412, including RGB (red, green, and blue). The display 412 can be any output device including a printer, a handheld device, and the like. The system 400 also includes a
memory 414 which can provide storage for the display 412, the color conversion block
25 410, and the bitstream conversion block 404. The memory 414 can be any type of storage such as dynamic random access memory (DRAM), extended output DRAM (EDO DRAM), synchronous DRAM (SDRAM), video ram (VRAM), static ram (SRAM), and the like. Alternatively, the bitstream conversion block 404, the color conversion block
410, and the display 412 can have their own local memory. Also, the decompression
30 block 408 can have its own storage or use the memory 414 for its storage when needed.

Fig. 8 illustrates a sequential decoding method 500 for decoding Huffman encoded data. The sequential decoding method 500 uses the following four tables to produce the decoded value:

$MINCODE[I]$ – minimum value of code words with length I ;
 $MAXCODE[I]$ – maximum value of code words with length I ;
 $HUFFVAL[J]$ – table of symbol values corresponding to the J^{th} code word;
 and

5 $VALPTR[I]$ – the index to the start of the list of values in $HUFFVAL$
 which are decoded by code words of length I .

The minimum length of all code words is $MinLength$ and the maximum length of all code words is $MaxLength$. Hence, the size of tables $MINCODE$,
 10 $MAXCODE$, and $VALPTR$ is equal to $[MaxLength-MinLength+1]$. The size of $HUFFVAL$ depends on the number of code words, and is denoted as N . The three tables $MINCODE$, $MAXCODE$, and $VALPTR$ are used to decode a pointer to the $HUFFVAL$ table for each valid Huffman code.

In a step 502, bitstream data is received for decoding. In a step 504, a
 15 current code length I is set to $MinLength$. Step 504 also initializes a variable $CODE$ to $NEXTBITS[MinLength]$ which contains the code words from the bitstream provided by the step 502. A step 506 compares the value of $CODE$ to $MAXCODE[I]$. If the value of $CODE$ is greater than the value of $MAXCODE[I]$, it is indicated that the current code word has a length larger than I , which was first initialized to $MinLength$ in step 504, and
 20 the step 506 is followed by a step 508. In the step 508, the value of I is incremented by one. The step 508 also reads the next bit of the bitstream into $CODE$. The step 508 updates the value of $CODE$ by performing a shift logic left (SLL) on current value of $CODE$ by one bit, and filling the least significant bit of $CODE$ with the next bit from the bitstream ($NEXTBIT$).

25 On the other hand, if in the step 506, it is determined that the value of $CODE$ is not greater than the value of $MAXCODE[I]$, it is indicated that the current code word has a length equal to or less than I , and the step 506 is followed by a step 510. The step 510 decodes the symbol value. The step 510 computes a code value pointer J and then uses J to compute the symbol value. The step 510 sets J to
 30 $\{VALPTR[I-MinLength]+CODE-MINCODE[I-MinLength]\}$. The step 510 then computes the symbol value by setting $VALUE$ to $HUFFVAL[J]$. Accordingly, the step 510 performs three table look-ups, including $VALPTR$, $MINCODE$, and $HUFFVAL$, to compute the symbol value. The calculated symbol value ($VALUE$) is then provided to a

step 512 to be output. The sequential decoding method 500 is repeated for the next code word.

The Huffman table used for encoding and decoding can be customized for each implementation. But, most coding standards provide a default Huffman table.

- 5 Table 1 below illustrates a recommended Huffman table for luminance DC difference in accordance with the JPEG standard. Further information regarding the JPEG standard can be found in "JPEG - Still Image Data Compression Standard," Appendix A. ISO DIS 10918-1, Requirements and Guidelines, pp. F-26, Van Nostrand Reinhold, 1993, by William B. Pennebaker and Joan L. Mitchell, which is incorporated herein by reference,
- 10 in its entirety and for all purposes.

In the JPEG standard, *MinLength* is 2 and *MaxLength* is 16. Code words are further sorted to be grouped in length and in ascending value order in each length group. The values in *MINCODE* and *MAXCODE* are signed 16 bit integers.

HuffValue	CodeLength	CodeWord	HuffCode
0	2	00	0
1	3	010	2
2	3	011	3
3	3	100	4
4	3	101	5
5	3	110	6
6	4	1110	14
7	5	11110	30
8	6	111110	62
9	7	1111110	126
10	8	11111110	254
11	9	111111110	510

15

Table 1

In Table 1, **CodeWord** represents the actual code words from the bitstream; **CodeLength** is the length of each code word; **HuffCode** is the value of each **CodeWord**; and **HuffValue** is the Huffman code for each symbol.

20

Taking the values from Table 1 as an example, with *MinLength* = 2 and *MaxLength* = 9, the table values of the three tables utilized by the sequential decoding method 500 are as follows:

$$MINCODE[8] = \{0, 2, 14, 30, 62, 126, 254, 510\};$$

$MAXCODE[8] = \{0, 6, 14, 30, 62, 126, 254, 510\}$; and
 $VALPTR[8] = \{0, 1, 6, 7, 8, 9, 10, 11\}$.

For example, considering an incoming bitstream "010," the step 504 will
 5 set I to 2 and $CODE$ to "01." The step 506 will return YES because "01" is greater than
 $MAXCODE[2-2]$ which is 0. In the step 508, I is incremented to 3 and $CODE$ is set to
 "010." The step 506 is repeated again and this time it will return NO because "010" is
 less than $MAXCODE[3-2]$ which is 6. Then, the step 510 will set J equal to 1. The step
 510 also looks up the symbol value for $HUFFVAL[1]$ and outputs this symbol value
 10 ($VALUE$) in the step 512.

Accordingly, the sequential method 500 repeatedly shifts the bitstream into
 $CODE$ and compares the current code with the maximum code of the same length. Once
 code length is known, the decoding of each value needs two additions and three table
 look-ups. This algorithm is sequential in nature because code lengths are checked
 15 sequentially. That is, before checking length I , there is no indication of whether the
 actual code length will be $I, I+1, I+2$, etc.

Fig. 9 illustrates a parallel decoding method 600 for decoding Huffman
 encoded data in accordance with an embodiment of the present invention. For the sake of
 example, an implementation of the parallel decoding method 600, such as discussed with
 20 respect to Fig. 6, will be described. Also, sample JPEG values discussed with respect to
 Fig. 8 and Table 1 are utilized in the sample embodiment. The parallel decoding method
 600 can, however, be performed on any processor capable of SIMD executions including
 Intel® MMX™ microprocessors, the Sun® UltraSPARC® with Visual Instructions Set,
 Advance Micro Device, Inc.'s® 3DNow!™, Intel® MMX™ with streaming SIMD
 25 extensions, and the like.

As detailed above, MMX™ technology can perform operations on a unit
 of 64 bits in the same instruction. The unit of 64 bits can also be treated as 8 bytes, 4
 words, 2 double words, or 1 quad word. For the example in accordance with values from
 Table 1, the 64 bits can be treated as four 16-bit registers (or 4 words) because the
 30 maximum code word length in Table 1 is 9. For the example embodiment of Fig. 9, the
 MMX™ instructions $PSRL$, $PSLL$, and $PSRA$ are utilized to shift the contents of selected
 variables. $PSRL$ denotes pack shift right logic. $PSLL$ denotes packed shift left logical
 and $PSRA$ denotes packed shift right arithmetic.

The parallel decoding method 600 receives the bitstream in a step 602. In a step 604, the bitstream is arranged in 4-bit portions (or quadbits) with little-endian format. This means if the original bitstream is of the form $b_0, b_1, b_2, b_3, \dots$ (where b_i stands for byte i), then the converted bitstream will be of form: $b_1, b_0, b_1, b_0, b_1, b_0, b_1, b_0, b_3, b_2, b_3, b_2, b_3, b_2, b_3, b_2, \dots$

This conversion can be implemented using 5.5 cycles for each 16-bit word. In a step 606, values of I , $MASK$, and $CODE4$ are initialized. As discussed above for the JPEG example, I is initialized to 2. The step 606 initializes $CODE4$ to the first set of 4x16-bit codes (or quadbits) from the bitstream. The parallel decoding method 600 uses $MASK$ to mask out the bits for a comparison operation in a step 608. The initial value for $MASK$ is set to "1100,0000,0000,0000,1110,0000,0000,0000,1111,0000,0000,0000,1111,1000,0000,0000 b" or "0xC000,E000,F000,F800." This is the bit mask for code lengths 2 through 5 which will be checked in the first iteration of the parallel decoding method 600.

In a step 608, $CODE4$ is masked off by the value of $MASK$. The step 608 also sets $MAXCODE4$ to $MAXCODE4_P[I-I+3]$ which represents values of $MAXCODE4_P$ for $I, I+1, I+2$, and $I+3$. Accordingly, $MAXCODE4$ will hold the maximum values for the four successive code lengths. Because the four values for $MAXCODE4$ will be loaded in one single instruction, the $MAXCODE$ table for the parallel decoding method 600 can be generated as follows:

20

$$I = 0$$

while I is less than $(MaxLength - MinLength + 2)$ do

$$J = (I/4) * 4 + 3 - I \% 4$$

$$MAXCODE_P[I] = (2^{(16 - MinLength - J)} * MAXCODE[J]) \div 1$$

25

$$I = I + 1$$

end while

30

where "/" denotes division by truncation and "%" denotes modulo operation. The values of $MAXCODE_P$ are unsigned 16-bit integers. The $MAXCODE$ table only has to be generated once for each Huffman table.

Taking the same example as in the sequential Huffman decoding algorithm discussed with respect to Fig. 8, the $MAXCODE$ table for the parallel algorithm will have the following values:

$$MAXCODE_xP[8] = \{61441, 57345, 49153, 1, 65281, 65025, 64513, 63489\}$$

For this example, the dimension of this table must be a multiple of 4.

5 Also, zero is assumed for those code lengths that do not exist in the Huffman code table.

In a step 610, the value of *MAXCODE4* is compared with *CODE4* which was initialized and masked in the steps 606 and 608. In the embodiments which use the MMX™ technology, since the MMX™ instruction PCMPEQ compares only signed words, this comparison operation can be performed by one unsigned subtraction with saturation and one comparison with zero. The values of *MAXCODE+1* can be stored in *MAXCODE4*. If *CODE4* is found to be greater than *MAXCODE4*, then in a step 612, *I* is incremented by 4 and the *MASK* is updated by shifting its value right arithmetically by 4 bits. The new value of *MASK* will create the bit mask for the next remaining 4 code word lengths (6 through 9). After the step 612, the steps 608 and 610 are repeated.

15 Accordingly, for the example of Table 1, the steps 608 and 610 are repeated twice, once for code word lengths 2 through 5 and once for code word lengths 6 through 9.

Once the step 610 returns a YES, a step 614 determines which one of the four lengths being checked contains the symbol. Accordingly, the step 614 computes a value of *DIF* which is set to [*MAXCODE4-CODE4-1*] shifted logically right by 4 bits.

20 The step 610 also performs a table look-up to provide the found symbol value (*VALUE*). The step 614 looks up the symbol value (*VALUE*) in a table *HUFFVAL_P*.

For this example, assuming $1 < I < 17$, a *HUFFVAL_P* table can be created using $[I+16*(MAXCODE[I]-CODE)]$ as index. Let:

25 $M = \max(MAXCODE[I]-MINCODE[I]);$ and
 $I = 0, 1, 2, \dots, MaxLength-MinLength.$

The *HUFFVAL_P* table will have a size of $[16*(M+2)]$. Let *COENUM[I]* denote the number of code words whose code length is $(I+MinLength)$. Let *CODE[J]* denotes the code value of the J^{th} code with code length *I*. The entries of the large *HUFFVAL* table can be computed as follows:

$I = 0$

while *I* is less than $(MaxLength-MinLength + 1)$ do

```

        J = Q
        while J is less than CODENUM[I] do
            HUFFVAL_P[I+16*(MAXCODE[I]+1-CODE[J])] =
                HUFFVAL[VALPTR[I]+CODE[J]-MINCODE[I]]
5          J = J + 1
        end while
        I = I + 1
    end while

```

10 Accordingly, values in the *HUFFVAL_P* table will be the maximum size of unsigned 16-bit integers depending on the number of the Huffman codes. By combining the three table look-ups of the sequential method into one, the instructions needed to index a Huffman code value is significantly reduced. At the same time, reducing the number of table look-ups saves valuable register space. Also, note that

15 (*MAXCODE[I]+1-CODE*) was computed when determining the code length and is available for this table look-up.

In a step 616, the calculated symbol value (*VALUE*) is output. Once the code length of current Huffman code is determined in the step 614, the current 4x16-bit code word can be left shifted and the same number of bits from the next 4x16-bit code word can be shifted in. The decoding of next Huffman code can then be started.

20

Accordingly, the parallel decoding method 600 reads in 4 bits of data at a time from the bitstream. Each consecutive 4-bit lengths are checked in parallel to see if the current code word falls in that range. Instead of the required 16 checks to cover the range 1 through 16 in the sequential algorithm, only 4 checks are needed to cover the

25 same range. Another significant improvement is to combine the three table look-ups for decoding into one single table look-up. This has the effect of minimizing computations as well as register usage, which can be very valuable in many situations.

SAMPLE EMBODIMENT FOR LENGTHS 1 THROUGH 8

The example embodiment discussed above assumes that the code words

30 have lengths between 2 and 9. If, however, the code words have lengths between 1 and 8, the decoding can be done on units of 8 bits. A maximum of one check instead of 8 is needed for each code word. This requires that the bitstream be expanded into 8x8 bits. For example, if the input bitstream is "b0,b1,..." (where *bi* stands for byte *i*), the converted bitstream will be "b0,b0,b0,b0,b0,b0,b0,b0,b1,b1,b1,b1,b1,b1,b1,b1,..." The

initial bit mask will now be "10000000,11000000,11100000,
11110000,11111000,11111100,11111110,11111111 b" or "0x80C0,E0F8,F8FC,FEFF."

The table entries will be generated as follows:

```

5          I = 0
          while I is less than (MaxLength-MinLength + 1) do
              J = (I/8)*8 + 7 - I%8
              MAXCODE_P[I] = (2(8-MinLength-I)*MAXCODE[J]) + 1
              I = I + 1
10         end while

          I = 1
          while I is less than 9 do
              J = 0
15         while J is less than CODENUM[I] do
              HUFFVAL_P[I+8*(MAXCODE[I]+1-CODE[J])] =
                  HUFFVAL[VALPTR[I]+CODE[J]-MINCODE[I]]
              J = J + 1
          end while
20         I = I + 1
        end while

```

The values for *HUFFVAL_P* table will be unsigned 8-bit integers. The values for *MAXCODE_P* table will be unsigned 8-bit integers.

25 SAMPLE EMBODIMENT FOR LENGTHS 1 THROUGH 32

If the length of code words are between 1 and 32, the decoding can be done on units of 32 bits. A maximum of 8 checks instead of 16 checks will be needed to determine the length of a code word. The bitstream will be extended into 2x32 bit form as follows:

```

30      Input bitstream: b0,b1,b2,b3,b4,b5,b6,b7,...
      Output bitstream: b3,b2,b1,b0,b3,b2,b1,b0,b7,b6,b5,b4,b7,b6,b5,b4,...
      The initial bit mask will now be "10000000,00000000,00000000,
      00000000,11000000,00000000,00000000,00000000,00000000 b" or "0x8000,0000,
      C000,0000."

```

The table entries will also be computed as follows:

```

I=0
while I is less than 32 do
5      J = (I/2)*2 + 1 - I%2
      MAXCODE_P[I] = (2(32 - MinLength-I)*MAXCODE[J]) + 1
      I = I + 1
end while

10      I = 1
      while I is less than 32 do
          J = 0
          while J is less than CODENUM[I] do
              HUFFVAL_P[I+32*(MAXCODE[I]+1-CODE[J])] =
15                  HUFFVAL[VALPTR[I]+CODE[J]-MINCODE[I]]
              J = J + 1
          end while
          I = I + 1
      end while
20

```

Thus, the values for *MAXCODE_P* table will be unsigned 32-bit integers. Values of table *HUFFVAL_P* will be the maximum size of unsigned 32-bit integers.

PERFORMANCE EVALUATION

25 The proposed parallel Huffman decoding algorithm has been implemented in MMX™ assembly code. For this section, the CPU time to expand the original bitstream into 4x16-bit portions is included in the decoding time.

If code words of length *I* appear with a probability of 2^{-I} , then the probability that a code word is of length *I* will be $CODENUM[I]*2^{-I}$. Letting *T[I]* be the
30 time to decode a code word of length *I*, then the average decoding time for one symbol can be calculated as follows:

$$\bar{T} = \sum_{I=2}^{16} \frac{CODENUM[I] * T[I]}{2^I}$$

Assuming the compression of ten common intermediate format (CIF) pictures using a sequential algorithm and where only one-fourth of the discrete cosine transform (DCT) coefficients are coded using the RUN-LEVEL symbols, the number of Huffman code words to be decoded can be calculated as follows:

$$(352*288/4)*10 = 253,440$$

Even though DCT is selected for this example, it would be obvious to those with ordinary skill in the art to utilize any transforms such as differential parse code modulation (DPCM), wavelet, sub-band transform, vector quantization, and the like. Using a 166 MHz Intel® Pentium® processor with MMX™, the decompression timing for the best case (where all code words are 2 through 5 bits long), worst case (where all code words are 12 through 15 bits long), and average case (where code lengths are distributed between 2 through 16 bits) are shown in Table 2 below.

Decoding Time (in msec)	Parallel Huffman Decoding
best case	111
average case	113
worst case	218

Table 2

Table 2 demonstrates that the parallel algorithm significantly reduces the worst case timing while giving superior performance for the average case. Also, note that even better results may be achieved by optimizing the assembly code further.

As will be understood by those with ordinary skill in the art, the present invention may be embodied in other specific forms without departing from the spirit or essential characteristics thereof. For example, the techniques of the present invention can also be applied to other SIMD systems such as the Sun® UltraSPARC® with Visual Instructions Set, Advance Micro Device, Inc.'s® 3DNow!™, and the like can be utilized.

Also, Intel®'s new Pentium® III (also know as Katmai), the successor of Pentium® II adds streaming SIMD extensions. Pentium® III provides 70 new instructions. Many instructions are for floating point SIMD operations which are targeted at 3-D applications. Several instructions for fixed point number operations have also been

added. These instructions are, in part, targeted at the MPEG encoding/decoding (motion estimation, motion compensation). Those with ordinary skill in the art would, however, understand that these new instructions can be taken advantage of to further accelerate the decoding techniques of the present invention. For example, Min and Max operations can speed up the code length detection. Insertion instruction to transfer integer registers and part of MMX™ registers also provides speed-up. Similarly, broadcast instruction can duplicate the bitstream in a more efficient manner.

Alternately, the techniques of the present invention can be implemented in a computer system. For example, the methods 500 and 600 can be implemented on a peripheral component interconnect (PCI) card. The PCI card can be installed onto the PCI bus of a personal computer. Also, other bus technologies such as NUBUS, ISA, EISA, Universal Serial Bus (USB), 1394 Firewire, and Accelerated Graphics Port (AGP) can also be utilized. Moreover, the techniques of the present invention can be implemented by utilizing the available routines and features such as caching, new instruction sets, multi processor systems, and their equivalents.

III. VIGNETTING CORRECTION.

Register 34 of Fig. 2 allows a vignetting correction to be applied, thus allowing a cheaper lens 10 to be used in camera 32. This allows the cost of the camera to be driven even lower.

In operation, the processor in host 36 will interrogate the camera to read the value of register 34. This will then be used in the vignetting correction block 40 of the host to correct for vignetting defects of the camera. This is done prior to pixel correction of block 42, to avoid vignetting being mistaken as a defective pixel. The correction can be accomplished by adding an offset value to the pixel brightness value provided.

Fig. 10 illustrates a pixel array 11 and a corresponding brightness curve 13 over the Y axis. A central pixel 17 with value (x_0, y_0) corresponds to the point of maximum brightness of the lens. A second example pixel 15 is near the minimum brightness, showing a vignetting effect of having the difference between curve 13 and a maximum brightness value 19. The present invention corrects this vignetting effect in both X and Y directions.

In one embodiment, a threshold indicated by r_T is illustrated within which the curve 13 is near the maximum value 19, and no vignetting correction need be applied.

Alternately, the correction can be applied to all of the pixels regardless to avoid any transition effects at the points r_T .

Depending upon the particular lens which is chosen, its specified vignetting amount can be used to provide a constant in a register 72. Thus, different lenses could be used in manufacturing a camera by simply programming a different constant value into register 72. This allows multiple sources to be used for the lenses, and allows compensation for changes in lens quality from lot to lot or manufacturer to manufacturer.

An embodiment of an equation for implementation and hardware or software to perform the vignetting correction is set forth below. Although the constant is stored in the camera, the vignetting correction can be done by either hardware/software in the camera, or software in a host. The vignetting defect of a low quality lens may be modeled as pass through in the middle of the lens and a parabolic curve near the edge of the lens. The transform function is isotropic. Let p_i and p_o be the pixel value before and after vignetting correction, and x_0 and y_0 be the coordinate of the center of the lens (should be the center of the image if the lens is mounted properly). Then,

$$\begin{aligned} p_o &= p_i * (a * r^2 + 1) \\ &= p_i * (a * ((x - x_0)^2 + (y - y_0)^2) + 1) \end{aligned}$$

"a" is a small positive constant. To be on the safe side, a maximum of 50% vignetting can be corrected, which means:

$$\begin{aligned} 1.0 &= 0.5 * (a * (176 * 176 + 144 * 144) + 1) \\ a &= 1.9338 * 10^{-5} \end{aligned}$$

Normalized by 2^{26} , this will give $a = 1298$. By limiting a to 0 through 1023, and performing the inverse computation, the maximum correctable vignetting is 56%.

Note that in the formula given above, the pixel aspect ratio is not taken into account. For square pixel, the equation is exact, for 12/11 pixel aspect ratio, the result is slightly off. Also, a will be scaled by $1/4$ if the sensor format is non-scaled QCIF (subsampling QCIF from CIF by leaving out every other pixels in both horizontal and vertical directions).

If y_i is 10 bits, $(x - x_0)^2$ and $(y - y_0)^2$ is 16 bits (for CIF size), and y_0 is 10 bits, this determines that a needs to be 10 bit (0 through 1023).

Fig. 11 is a block diagram of one embodiment of a hardware system or software blocks for implementing the equation described above. In particular, an optional selection circuit/block 80 is applied to the select input of a multiplexer 82. Where the pixel position is less than the threshold r_2 , the input pixel p_i on line 84 is simply passed through to the output pixel, p_0 , on line 86. If it is greater than the threshold, a vignetting correction is applied by implementing the formula above with the hardware shown in Fig. 11.

In particular, the current x value of the pixel, x_i , is applied to arithmetic circuit/block 88, where the central pixel, x_0 , is subtracted from it. This value is then squared by applying its input twice, through a multiplexer 90, to a multiply circuit/block 92, where it is accumulated in an add/accumulator 94. Similarly, the y value y_i is provided to a circuit/block 96, where the difference from the center y value, y_0 , is determined. This is similarly applied as two values through mux 90 to a multiplier 92, where it is squared and provided to add/accumulator 94, where it is added to the squared x value. The use of the multiplexer simply allows, by time division multiplexing, the same multiplier 92 to be used for both calculations. Obviously, an alternative would be to provide two separate multiply units.

The output of accumulator 94 is then provided to a second multiplier 96, where it is multiplied by the constant value from register 72. The value of one is then subtracted in a unit 98, and this value is multiplied by the pixel value, p_i , in a multiplier 100. The output is provided on a line 102 through multiplexer 82 to output line 86. Obviously, variations of the circuitry could be provided, such as using the same multiply circuit for multipliers 92, 96 and 100, with the multiplier being used for different purposes in different time slots.

Optional circuit/block 80 provides a comparator 104 which compares the output of add/accumulator 94 (the radius value of the x^2+y^2 values) to the threshold radius in a register 106.

Preferably, the constant a has the same number of bits as the pixel value, p_i . Thus, if p_i is 10 bits, while the x and y values are 16 bits, the constant a would be 10 bits. Also, a can be scaled by $1/4$ if the sensor format is non-scaled QCIF (cropped). Thus, correction is provided where the image is cropped before the vignetting correction.

In addition, correction can be done for the pixel aspect ratio by multiplying the y value by $11/12$ where a TV will be used.

In one embodiment, the correction will take place only outside a predefined diameter. This assumes that the lens performs satisfactorily around the center of the lens, which is often a reasonable assumption. By doing so, a more accurate correction model can be derived. Let r_T be the diameter within which the vignetting defect can be ignored. Then,

If $r \leq r_T$,

$$p_o = p_i,$$

if $r > r_T$,

$$p_o = p_i * (a * (r^2 - r_T^2) + 1)$$

$$= p_i * (a * (((x - x_0)^2 + (y - y_0)^2) - ((x_T - x_0)^2 + (y_T - y_0)^2)) + 1)$$

The present invention thus allows variation in lenses and also cheaper lenses to be used. It improves the AGC and AWB in the host computer. Additionally, it allows better object motion tracking. Object motion is usually done by assuming the luminance value is the same as the object moves to the edge of a sensor, and thus tracking by looking for the same luminance value. Obviously, vignetting effects can thwart the efforts to locate the object. By applying the vignetting correction close to the sensor, this can be overcome. The invention also provides overall improved video quality.

IV. PIXEL CORRECTION.

Block 42 of Fig. 2 performs pixel correction in the host. This allows for defective detector locations on the CMOS or CCD sensor array to be corrected. In general, this is accomplished by comparing a brightness value to the brightness value of the neighboring detector elements. If the difference is more than a threshold value, it is assumed that this is due to a defective sensor element. Accordingly, a corrected value will be substituted, typically an average of the surrounding pixel values.

In order for this to be done in the host without affecting the frame rate of the video, the algorithm is kept simple. This is accomplished through two primary features. First, the defective pixel detection is not done on any frame, but only on a subsample of the frames. For example, the detection may be done only every 32-128 frames. Second, the statistics kept are simple, to reduce the amount of processing required. In particular, the statistics may include simply the location of the pixel element, and the frequency or number of times a defective pixel has been detected.

The system works by reading off or scanning the raw image data from an image pickup device for pixels that vary more than a specific amount in intensity from their neighboring pixels. The raw image sensor data is the unprocessed brightness data output obtained from the image sensor and which has not gone through any lossy compression or color processing. The image sensor reads analog voltage or current, converts to digital and sends the signal to a host with no further processing or compression. The photosites on the image sensor are used to capture either color or monochrome digital still or video images. The raw image data is sent to the intelligent host over a bus with a data transfer rate which is determined by the bus protocol of the particular bus such as a universal serial bus (USB) or a parallel port.

The raw image sensor data, the location and frequency of occurrence of each defective pixel, and the results of all intermediate computations performed by the computer program are all stored in memory. Other algorithms are then used to average the values of an anomalous pixel's neighboring pixels to replace the data from the defective pixel. The corrected data can then be further processed and ultimately displayed on the monitor. The process includes video subsampling, meaning that the detection is carried out and repeated at various frame intervals. The video subsampling is carried out on one of every 128 (1/128) frames. Alternately, the video subsampling can be carried out on every 1/64 video frames. And yet in an additional embodiment, the video subsampling is carried out on every $1/(n \text{ times } X)$ frames, where n is an integer and X is not equal to 50 or 60. 50 and 60 correspond to 50 Hz and 60 Hz, which are AC lighting frequencies used in the United States and Europe respectively. This way, it is ensured that anomalous raw data pixels are not artifacts of the artificial lighting systems. The use of video subsampling allows for rapid and optimum corrections without the need to scan every frame which would adversely impact the processing speed of the processor and the system. Naturally, no video subsampling is employed when detecting and correcting defective pixels in still images.

A statistical database recording the location and the frequency of occurrence of defective pixels is generated and stored in memory to allow the system to learn and adapt its operation over time. The trends from the statistical database are stored so that truly anomalous pixels can over time be distinguished from false detection of true anomalies in the target image, lighting or other environmentally induced anomalies. The operations logic of the statistical database and the details of the operation of the computer program are described below.

Fig. 12 illustrates a simplified flow chart describing the functioning of the software program implemented in a system for detecting and correcting defective pixels according to one embodiment of the present invention. The first step in defective pixel detection and correction includes acquiring a portion of a frame of raw image data from an image sensor array, step 1210. The raw image data may be data corresponding to a live scene being digitized or it may be data corresponding to a calibration or "black" background image. The "black" background image may be obtained by reducing the integration time of the sensor array or by reducing the gain of the image signal. The raw data from the "black" image can be used to detect over active photosites. The raw data corresponding to the "black" image must also be black, and if any pixel is not, then it corresponds to an overactive photosite. The use of the "black" image can enhance defective pixel detection by removing any brightness deviations that originate from the scene itself.

Although the data from an entire frame can be processed at one time for defective pixel detection and correction, the algorithm begins its function as soon as data from three lines from the image sensor has arrived at the host. Three lines of data from the image sensor array will allow for the processing of data for a pixel and all its surrounding neighbors. Moreover, when the sensor array is being used to capture video images, video subsampling is used so that not every frame of the video captured by the image sensor array is required for the defective pixel detection operation. In a preferred embodiment, defect detection is carried out at predetermined frame intervals as described above and the correction is applied to all video frames. The judicious use of video subsampling allows for the method to be very fast and efficient.

In an alternate embodiment, no video subsampling is used such that the defective pixel detection is carried out on every video frame. The use of video subsampling is a function of the intelligent host's processor speed. If a processor is fast enough to allow detection of anomalous pixels in every frame, then video subsampling is not used. If, on the other hand, the processor is not fast enough to allow video processing at a desired frame rate, then video subsampling is used to ensure data transfer at that desired frame rate. Without video subsampling, anomalous pixel correction is immediate, such that defects are corrected in the very frame in which they are detected. With video subsampling, anomalous pixel correction is delayed until a frame is sampled to detect anomalous pixels. Therefore, the choice of using video subsampling, and the rate of

subsampling are a function of the processor speed and a trade off between processor power and the delay before correction.

After data acquisition, for each of the acquired pixels, the local pixel brightness value and also the average brightness value of all its immediately neighboring pixels are computed and stored, step 1212. Next, a deviation threshold value is established, step 1214. The deviation threshold establishes the acceptable level of variance between a pixel's brightness value and the average brightness value of all its immediately neighboring pixels. Next, for each acquired pixel, a local brightness deviation is computed, step 1216. The local brightness deviation is the absolute value of the difference between a pixel's brightness value and the average of the brightness value of all its immediately neighboring pixels.

Next, for each pixel whose data that has been acquired, its local brightness deviation is compared to the deviation threshold, step 1218. Any pixel whose local brightness deviation exceeds the threshold deviation value is then flagged as a defective pixel. The physical location and the frequency of occurrence of each defective pixel is then recorded in a statistical database, step 1220. The statistical database is then queried, to determine whether the defective pixel's data value should be corrected, step 1222. The statistical database, by storing the location and frequency of defective pixels, develops over time trends which confirm which of the defective pixels warranted correction. The logic of the trends from the statistical database initially warrant correction of all flagged defective pixels as a default, and over time warrant pixel correction only if a particular pixel has an occurrence frequency of at least two out of the last four queries.

The defective pixels that have passed through the statistical database filter are corrected next, step 1224. In order to correct a defective pixel, the erroneous raw brightness data for a defective pixel is replaced by that pixel's local average brightness value, which is the average brightness value of all its immediately neighboring pixels. Next, the corrected data from the defective pixels as well as data from non-defective pixels is prepared to be sent for subsequent processing, step 1226. Subsequent processing may include compression, color processing and encoding to data formats suitable for display. The defective pixel detection and correction is carried out on the raw data from the image sensor array because it is preferred to correct the data before any subsequent processing has occurred since that processing itself can introduce artifacts which are hard to distinguish from artifacts which have been produced as a result of defective photosites.

The software algorithm is intentionally kept very simple so that the processor load and/or the video frame rate is minimally affected. First, the software algorithm is kept simple because it only performs three functions, namely the detection, correction and statistics functions. The statistics routine is kept simple because only
5 position information and frequency of occurrence of anomalous pixels are tracked. Second, the software algorithm is kept simple so as to have a minimal impact on the rate at which data is transferred to the host, so that while a frame is being scanned for defective pixels, subsequent frames are not held up. The maximum impact of the software algorithm is to at worst reduce the video data transfer rate from 10 frames per
10 second (fps) to 9 fps. Third, the software algorithm is kept simple such that the host processor's load is not increased by more 1 % when executing the algorithm. The anomalous pixels are detected in one frame and the corrections are later instituted on subsequent frames as they are read off the image sensor.

Therefore, when the system is operating, as raw image data from the image
15 sensor is obtained and sent to be processed on a host, frames are captured for defective pixel detection and correction as described above. Once a defective pixel has been identified and its correction has been warranted by the statistical database, that correction will continue for all subsequent frames until the next query is made to the statistical database. If at that time the initially flagged defective pixel continues to be designated as
20 a defective one, then the corrections will continue as described above. If on the other hand, the statistical database does not warrant the correction of defective pixels, then those pixels will no longer be corrected until the next query to the statistical database is made. The dynamic nature of this system will also allow for corrections of data from photosites that become defective over time due to post-manufacture environmental
25 changes.

The statistical analysis segment of the defect detection and correction algorithm is an optional one. It is aimed at increasing the efficiency of the pixel corrections, so as not to correct anomalies that were not caused by defective photosites, and hence save processing time and load. However, the gains in efficiency must be
30 balanced against the load imposed by the statistical analysis portion itself. In an alternate embodiment not employing the statistical analysis portion, all pixels that get flagged as defective get corrected. As in the choice of using or not using the video subsampling, the decision to employ the statistical analysis portion of the defect detection and correction algorithm depends on a trade off between efficiency and processor power. If the

processor is fast enough, then efficiency concerns are not so important. On the other hand, if the processor is not fast enough, then pixel correction efficiency becomes important enough to warrant the implementation of the statistical analysis portion of the defect detection and correction algorithm.

5 As will be understood by those of skill in the art, the present invention may be embodied in other specific forms without departing from the essential characteristics thereof. For example, the rate of video subsampling could be varied, or not done at all. Alternately, the logic of the optional statistical database could be altered from one of correcting defective pixels as a default to one where defective pixels are
10 corrected only if warranted by the trends from the statistical database.

V. NO FRAME BUFFER MEMORY.

 The elimination of the processing in the camera allows the elimination of the frame buffer memory, since digital pixel data need not be stored for such processing.
15 However, another potential problem is that the output of data from the sensor needs to be held pending availability of the bus. This can require a very large buffer memory. In one embodiment of the camera, only a small buffer in the bus interface is used. A small buffer is achievable by controlling the timing of when the sensor is read to correspond to when the host bus is available, and reading only an amount of data which can be buffered
20 until the next bus availability.

 Fig. 13 is a block diagram of a video camera according to the present invention. This system includes a lens 1310 and a sensor array, such as a CMOS sensor array 1312, and is connected to a USB 1320.

 The digitized signals from the CMOS sensor chip, rather than being stored
25 in a frame buffer memory as in the prior art, are provided directly to a digital processing circuit 1348. They are first provided to a video digital signal processor 1350 which performs the compression, and optionally order functions to minimized data transfer, such as cropping, scaling and digital filtering. Once processed, the digital data is provided to a bus interface 1354.

30 Bus interface 1354 includes a bus controller buffer 1356 and a bus controller 1358. Buffer 1356 stores at least a single line of data from a CMOS sensor row. Preferably, bus controller buffer 1356 is capable of storing two USB frames at the maximum possible rate, or 2 Kbytes. Depending upon the frame rate, compression rate, etc., this may typically vary up to 10 lines, or even to 15-20 lines if decimation is done.

Bus controller 1358 provides a control signal on a line 1360 to timing and control generator 1326. Timing generator 1326 provides clocking signals on line 1364 to CMOS sensor array 1322. Clocking signals 1364 include the row and column transfer pulses. The column transfer pulses are provided as in the prior art, periodically loading the charged values from the CMOS sensor array. However, the row transfer pulses are varied to match the reading of the data out of the CMOS sensor array with the processing by the following circuitry and the transferring of data to the USB.

In one embodiment, the digital processing circuit 1348 is integrated onto the same semiconductor chip substrate as CMOS sensor array 1312, timing generator 1326, ADC 1328, and the row and column logic and buffers. The elimination of the frame buffer allows this single chip. This allows a more compact, less expensive video camera to be built.

Fig. 14 illustrates the timing of the row transfer pulses provided on line 1364. Rather than a continuous stream of pulses, the bits are clocked out as needed by the processing circuit in groups as illustrated. A first group of three lines 1466 is issued, and then there is a delay while these pixels are processed and transferred to the bus. Then, a next of group of two lines 1468 may be provided. There might then be a delay depending upon the availability in the line buffer and the timing of transfer of data onto the USB. A single line 1469 is shown being transferred next, followed by a subsequent group of four lines transferred by a number of pulses 1470. After an entire image has been transferred, a vertical transfer pulse 1472 is provided. The example of Fig. 14 is intended to be simply illustrative. Note that because the row transfer pulses are provided between the same column pulses, in order to allow the gaps shown, the pulses are provided much closer together. Thus, bursts of small groups of pulses are provided to quickly send pixel data as fast as the circuitry can handle it. In one embodiment, an entire image may be clocked out in a series of closely spaced pulses, allowing the processing to complete before the next column transfer pulse. If the processing still is not completed at the time of the next column transfer pulse, there can be a delay, which conforms to the timing of the processing and the availability of the USB. Thus, the row pulses may all be sequentially provided, but at a faster rate than in the prior art and the timing of this grouping of pulses after the column pulse can be varied to be right after, or just before the next column pulse, in order to accommodate the processing and bus timings.

VI. NO COMPRESSION.

In one embodiment of the invention, the compression block 26 and decompression block 38 of Fig. 2 can be eliminated. This may be done, for example, when a new USB is developed with higher bandwidth. However, prior embodiments may still be desirable if increased sensitivity sensors are developed, allowing higher resolution video image data to be transmitted, using up the additional bandwidth made available. However, one embodiment takes advantage of the additional bandwidth by eliminating the need for compression, thus further simplifying the video camera and reducing its cost.

The same major functions shown in Fig. 2 would be performed in the host, with the exception of the decompression. The camera would still include the correction register 34 in one embodiment, allowing for the inexpensive lens 10 to be used.

VII. LOSSY COMPRESSION.

In an alternate embodiment, a lossy compression method may be used. Lossy compression means that the recovered image may lose some resolution. A disadvantage of lossy compression is that any vignetting or pixel correction done in the host would be degraded. Accordingly, in this alternate embodiment, the vignetting correction and the pixel correction is done in the video camera itself. This allows a lossy compression block to be used for block 26 of Fig. 2. The color processing and scaling operations are still performed in the host, thus providing a simpler camera than the prior art of Fig. 1.

As will be understood by those of skill in the art, the present invention may be embodied in other specific forms without departing from the essential characteristics thereof. For example, in one embodiment the correction register can correct for features of a camera that can vary from camera to camera on manufacture, other than features of the lens. Instead of a physical shared bus, the data can be wirelessly transmitted from the camera to the host, or to a receiver connected to the host. The host itself can be a personal computer, an intelligent Internet device, or any other appliance or component with some sort of processor or processing circuitry. Accordingly, the foregoing description is intended to be illustrative, but not limiting, of the scope of the invention which is set forth in the following claims.

WHAT IS CLAIMED IS:

1 1. A video camera system comprising:
2 (a) a video camera having
3 an image sensor, and
4 a processing element for performing data compression;
5 (b) an interface coupled to said video camera; and
6 (c) a processing unit for receiving compressed data from said interface,
7 said processing unit being configured to perform
8 decompression of said compressed data, at least a portion of the operations
9 of said decompression being performed in parallel,
10 anomalous pixel correction, and
11 color processing;
12 such that anomalous pixel correction and color processing circuitry is not
13 needed in said video camera, and said anomalous pixel correction is performed prior to
14 said color processing.

1 2. The video camera of claim 1 wherein said video camera includes a lens,
2 and said processing unit is further configured to perform vignetting correction for said
3 lens.

1 3. The video camera of claim 1 further comprising a memory element in
2 said video camera for storing a constant value corresponding to an amount of vignetting
3 of said lens.

1 4. The video camera of claim 1 wherein said anomalous pixel correction is
2 designed to minimize its impact on video frame rate by including video subsampling to
3 determine defective pixels and statistical recording of only pixel position and data related
4 to the frequency of detected defective pixels.

1 5. The system of claim 1 wherein said processing unit comprises a
2 processor having instructions for operating on multiple pixels in a register in parallel, said
3 processor being programmed for decompression a data stream of variable size
4 compression codes, wherein at least a portion of the operations of the decompression are
5 performed in parallel on multiple fields in said register.

1 6. The system of claim 5 wherein said data stream is compressed using
2 Huffman encoding, and said operations include:
3 duplicating a group of bits from said data stream to provide duplicates in
4 multiple positions in a first register; and
5 performing an operation between values in a second register and said
6 duplicates in parallel, with a different value being used for each duplicate.

1 7. A video camera system comprising:
2 (a) a video camera having
3 an image sensor, and
4 an analog to digital converter coupled to said image sensor;
5 (b) an interface coupled to said video camera, said interface being a shared
6 interface capable of concurrent use by other peripheral devices; and
7 (c) a processing unit for receiving digital data from said interface, said
8 processing unit being configured to perform
9 anomalous pixel correction, and
10 color processing;
11 such that anomalous pixel correction and color processing circuitry is not
12 needed in said video camera, and said anomalous pixel correction is performed prior to
13 said color processing.

1 8. The system of claim 7 wherein said video camera further includes a
2 processing device for compressing digital data from said analog to digital converter.

1 9. A method for decoding a Huffman compressed data stream in a
2 processor having instructions for operating on multiple pixels in a register in parallel,
3 comprising:
4 duplicating a group of bits from said data stream to provide duplicates in
5 multiple positions in a first register; and
6 performing an operation between values in a second register and said
7 duplicates in parallel, with a different value being used for each duplicate.

1 10. The method of claim 9 further comprising:
2 using a result of said operation in an indexing operation to index a table of
3 Huffman decoded values.

1 11. A video camera comprising:
2 an image sensor;
3 an analog to digital converter coupled to said image sensor;
4 a memory element for holding a value corresponding to a correction factor
5 for an aspect of said camera, said value being transferable to a remote host; and
6 an interface, coupled to said analog to digital converter and said memory
7 element, for establishing communication with said remote host.

1 12. The camera of claim 11 further comprising a lens, and wherein said
2 value is a lens vignetting value.

1 13. A method of detecting and correcting defective pixels in raw data
2 taken from an image sensor used to obtain a digitized image, wherein said raw data
3 includes normal pixels and defective pixels, said method comprising the steps of:
4 (a) receiving a raw data signal for each pixel in said image;
5 (b) computing for each pixel received from said image sensor a brightness
6 value;
7 (c) computing for each pixel received from said image sensor a local
8 brightness value;
9 (d) computing for each pixel received from said image sensor a local
10 brightness deviation of said brightness value from said local brightness value;
11 (e) setting a deviation threshold;
12 (f) comparing for each pixel received from said image sensor, its local
13 brightness deviation to said deviation threshold and designating pixels having local
14 brightness deviations greater than said deviation threshold as defective pixels;
15 (g) recording the location of said defective pixels in a statistical database;
16 (h) recording the frequency of occurrence of said defective pixels in said
17 statistical database; and
18 (i) correcting the brightness value of said defective pixels, provided said
19 correcting is warranted by trends from said statistical database.

1 14. A method for operating a video camera comprising the steps of:
2 reading image data from an image sensor array;

3 processing said image data in a video digital signal processor (DSP), said
4 processing including performing image compression;
5 providing processed image data to a bus; and
6 providing a clock signal to said image sensor array to transfer image data
7 from said image sensor array at a rate corresponding to a data transfer rate over said bus,
8 such that a frame buffer memory for storing an entire frame of said image is not needed.




FIG. 1. PRIOR ART




FIG. 2.




FIG. 3.

3/11




FIG. 4.




FIG. 5.




FIG. 6.


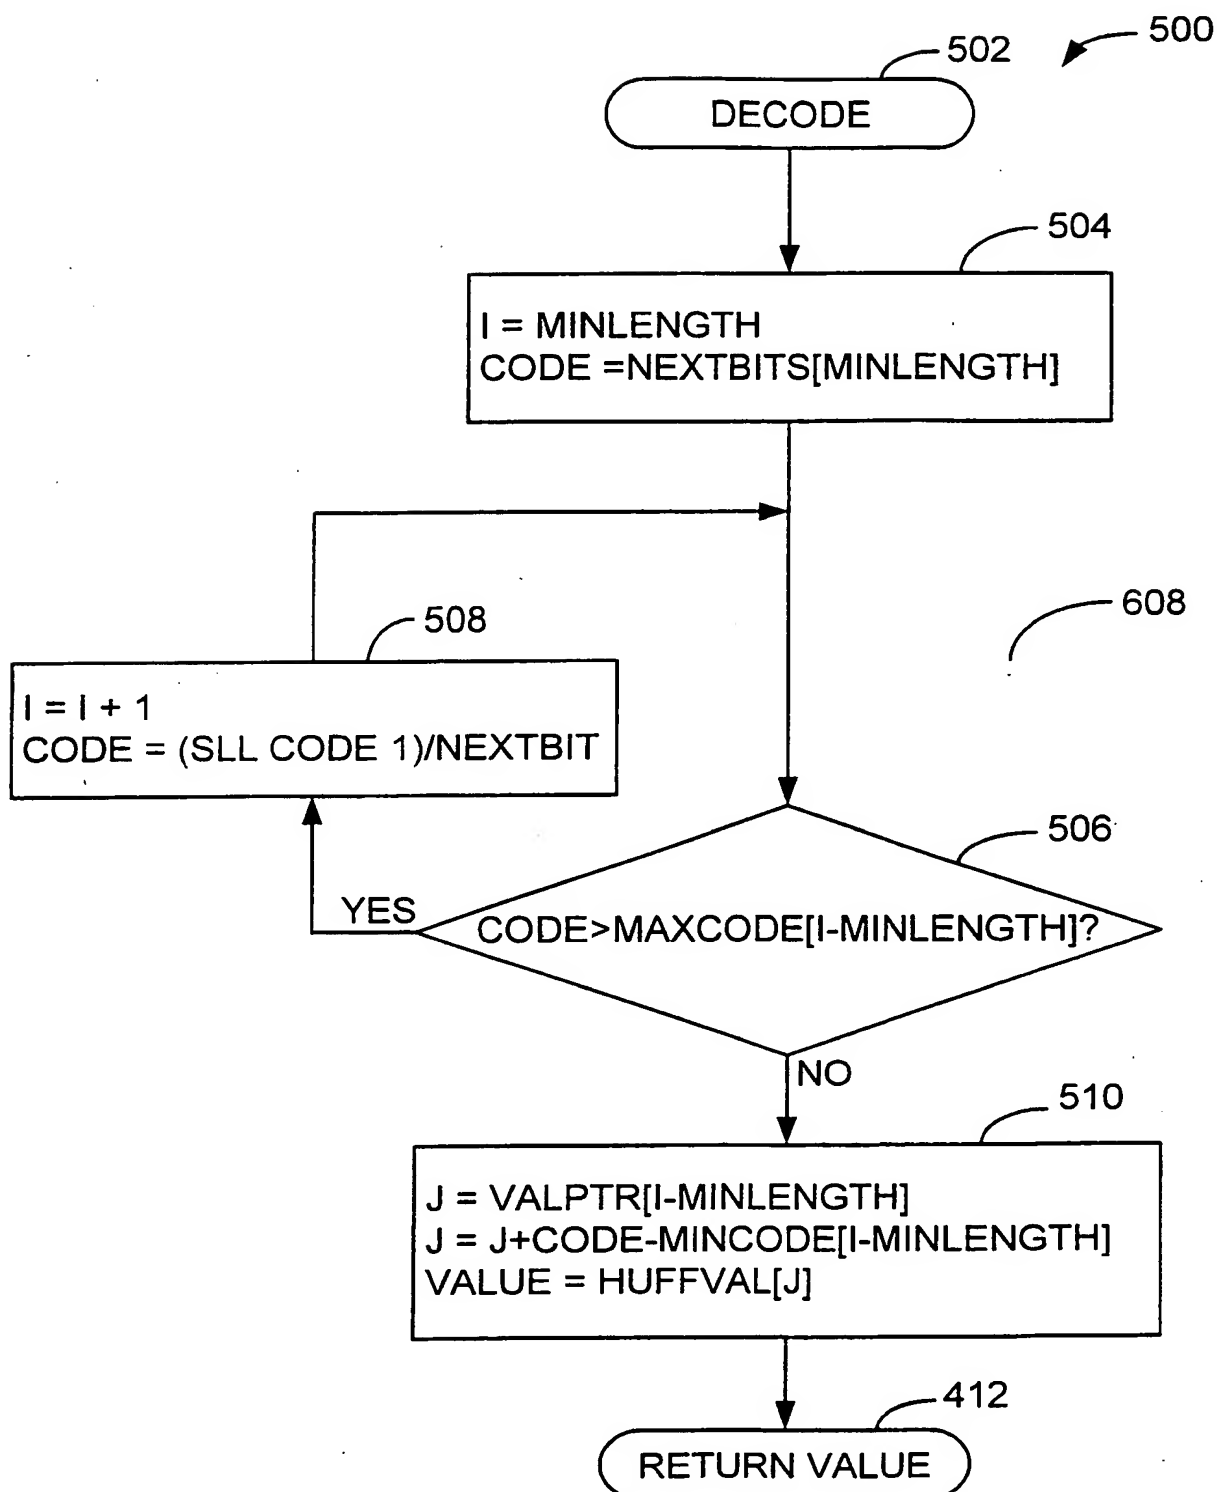
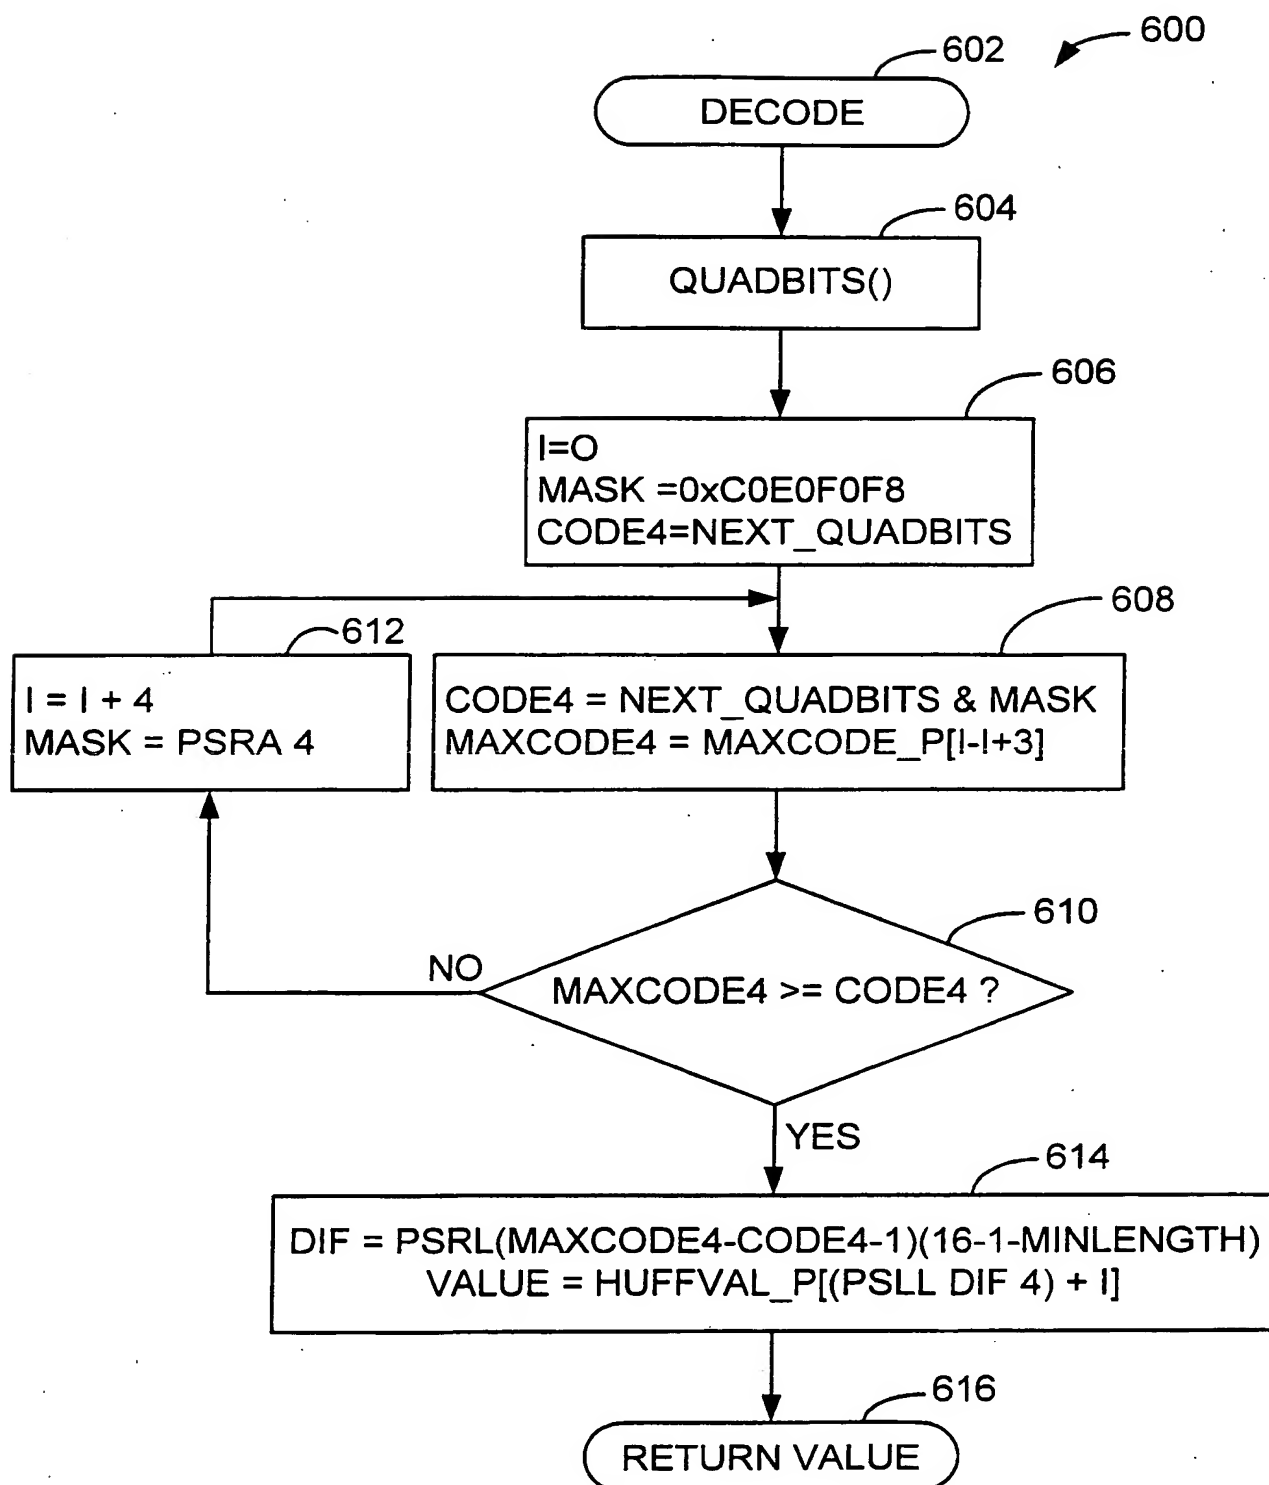


FIG. 7.

6/11

**FIG. 8.**

7/11

**FIG. 9.**

8/11




FIG. 10.




FIG. 11.

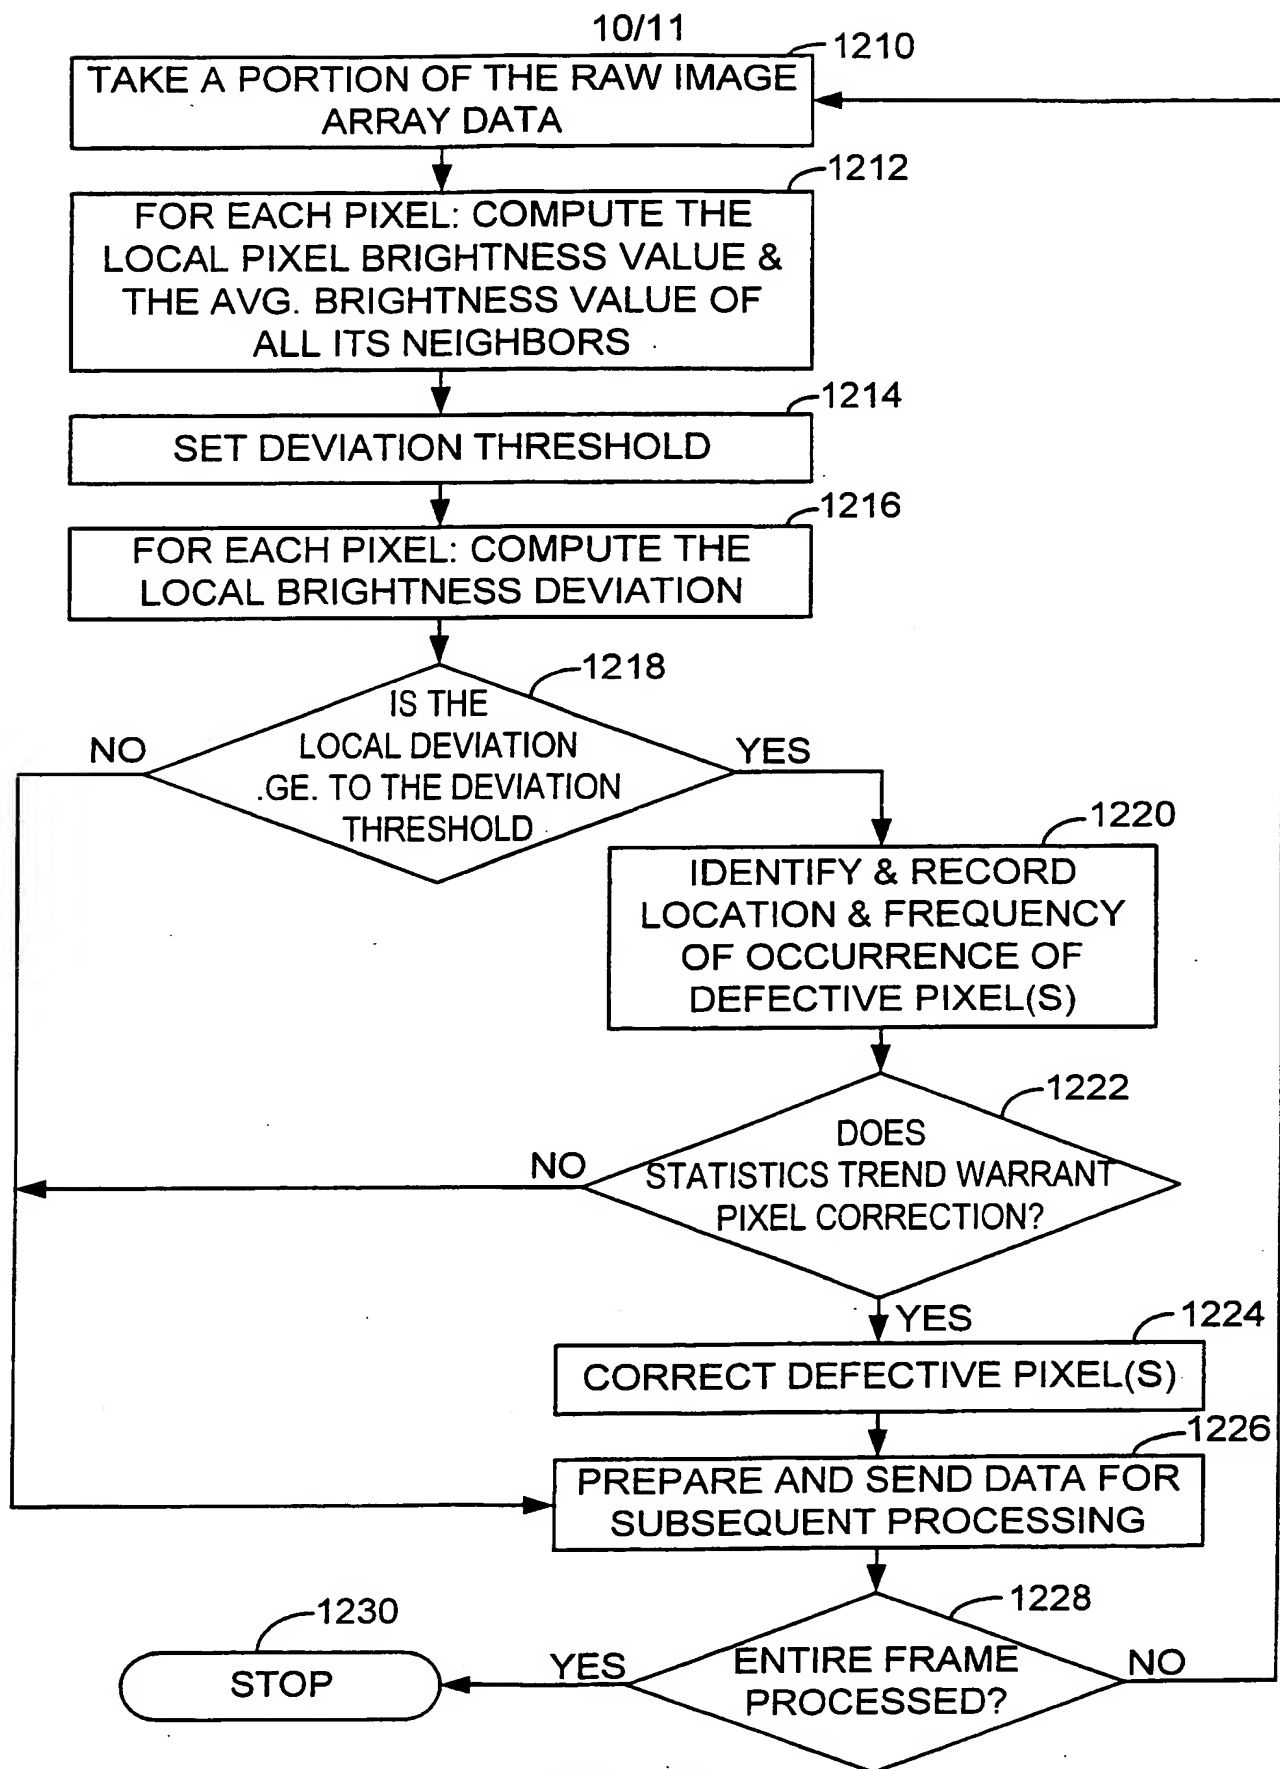


FIG. 12.
SUBSTITUTE SHEET (RULE 26)




FIG. 13.




FIG. 14.